arXiv:1407.4755v1 [cs.CC] 17 Jul 2014

On The Communication Complexity of Linear Algebraic Problems in the Message Passing Model

Yi Li* Xiaoming Sun† Chengu Wang‡ David P. Woodruff§

Abstract

We study the communication complexity of linear algebraic problems over finite fields in the multi-player message passing model, proving a number of tight lower bounds. Specifically, for a matrix which is distributed among a number of players, we consider the problem of determining its rank, of computing entries in its inverse, and of solving linear equations. We also consider related problems such as computing the generalized inner product of vectors held on different servers. We give a general framework for reducing these multi-player problems to their two-player counterparts, showing that the randomized s-player communication complexity of these problems is at least s times the randomized two-player communication complexity. Provided the problem has a certain amount of algebraic symmetry, which we formally define, we can show the hardest input distribution is a symmetric distribution, and therefore apply a recent multi-player lower bound technique of Phillips *et al.* Further, we give new two-player lower bounds for a number of these problems. In particular, our optimal lower bound for the two-player version of the matrix rank problem resolves an open question of Sun and Wang.

A common feature of our lower bounds is that they apply even to the special "threshold promise" versions of these problems, wherein the underlying quantity, e.g., rank, is promised to be one of just two values, one on each side of some critical threshold. These kinds of promise problems are commonplace in the literature on data streaming as sources of hardness for reductions giving space lower bounds.

*Max-Planck Institute for Informatics. yli@mpi-inf.mpg.de.

†Institute of Computing Technology, Chinese Academy of Sciences. sunxiaoming@ict.ac.cn.

‡Google Inc. wangchengu@gmail.com.

§IBM Almaden Research Center. dpwoodru@us.ibm.com.

1 Introduction

Communication complexity, introduced in the celebrated work of Yao [36], is a powerful abstraction that captures the essence of a host of problems in areas as disparate as data structures, decision trees, data streams, VLSI design, and circuit complexity [16]. It is concerned with problems (or *games*) where an input is distributed among $s \geq 2$ players who must jointly compute a function $f : X_1 \times \cdots \times X_s \to Z$, each X_i and Z being a finite set: Player i receives an input $x_i \in X_i$, the players then communicate by *passing messages* to one another using a predetermined protocol $\mathcal{P}$, and finally they converge on a shared output $\mathcal{P}(x_1, \ldots, x_s)$. The main goal of the players is to minimize the amount of communication, i.e., the total length of messages communicated. Put $\mathbf{x} = (x_1, \ldots, x_s)$. We say that a deterministic protocol $\mathcal{P}$ computes f if $\mathcal{P}(\mathbf{x}) = f(\mathbf{x})$ for all inputs $\mathbf{x}$. In a randomized protocol, the players can flip coins and send messages dependent on the outcomes; we shall focus on the public coin variant, wherein the coin flip outcomes are known to all players.[1] We say a randomized protocol $\mathcal{P}$ computes f with error δ if $\Pr[\mathcal{P}(\mathbf{x}) = f(\mathbf{x})] \geq 1 - \delta$ for all inputs $\mathbf{x}$. In all cases, we define the *cost* of $\mathcal{P}$ to be the maximum number of bits communicated by $\mathcal{P}$ over all inputs. We define the *deterministic* (resp. δ-error *randomized*) communication complexity of f, denoted $\mathrm{D}(f)$ (resp. $\mathrm{R}_\delta(f)$) to be the minimum cost of a protocol that computes f (with error δ in the randomized case). It holds that $\mathrm{D}(f) \geq \mathrm{R}_\delta(f)$ for all f and $0 \leq \delta \leq 1$.

Most work in communication complexity has focused on the two-player model (the players are named Alice and Bob in this case), which already admits a deep theory with many applications. However, one especially important class of applications is *data stream* computation [11, 22]: the input is a very long sequence that must be read in a few *streaming passes*, and the goal is to compute some function of the input while minimizing the memory (storage space) used by the algorithm. Several data stream lower bounds specifically call for *multi-player* communication lower bounds [2]. Moreover, several newer works have considered distributed computing problems with streamed inputs, such as the *distributed functional monitoring* problems of Cormode *et al.* [8]: in a typical scenario, a number of "sensors" must collectively monitor some state of their environment by efficiently communicating with a central "coordinator." Studying the complexity of problems in such models naturally leads one to questions about multi-player communication protocols.

In the multi-player setting, strong lower bounds in the message passing model[2] are a fairly recent achievement, even for basic problems. For the SetDisjointness problem, a cornerstone of communication complexity theory, two-player lower bounds were long known [14, 26] but an optimal multi-player lower bound was only obtained in the very recent work of Braverman *et al.* [5]. For computing bit-wise AND, OR, and XOR functions of vectors held by different parties, as well as other problems such as testing connectivity and computing coresets for approximating the geometric width of a pointset, optimal lower bounds were given in [25]. For computing a number of graph properties or exact statistics of databases, a recent work achieved optimal lower bounds [34]. There are also recent tight lower bounds for approximating frequency moments [33] and approximating distinct elements [35]. Our chief motivation is to further develop this growing theory, giving optimal lower bounds for other fundamental problems.

Linear algebra is a fundamental area in pure and applied mathematics, appearing ubiquitously in computational applications. The communication complexity of linear algebraic problems is therefore intrinsically interesting. The connection with data streaming adds further motivation, since linear algebraic problems are a major focus of data stream computation. Frieze, Kannan and Vempala [9] developed a fast algorithm for the low-rank approximation problem. Clarkson and Woodruff [7] gave near-optimal space bounds in the streaming model for many linear algebra problems, e.g., matrix multiplication, linear regression and low rank approximation. Muthukrishnan [23] asked several linear algebra questions in the

[1] Though the private coin model may appear more "natural," our key results, being lower bounds, are *stronger* for holding in the more general public coin model. In any case, for the particular problems we consider here, the private and public coin models are asymptotically equivalent by a theorem of Newman [24].

[2] In contrast to the message passing model is the *blackboard* model, where players write messages on a shared blackboard.

streaming model including rank-k approximation, matrix multiplication, matrix inverse, determinant, and eigenvalues. Sárlos [27] gave upper bounds for many approximation problems, including matrix multiplication, singular value decomposition and linear regression.

Our Results: Let us first describe the new two-player communication complexity results proved in this work. We then describe how to extend these to obtain our multi-player results.

Two-Player Lower Bounds: We start by studying the following closely related matrix problems. In each case, the input describes a matrix $z \in M_n(\mathbb{F}_p)$, the set of $n \times n$ matrices with entries in the finite field $\mathbb{F}_p$ for some prime p.

- Problem $\textsc{Rank}_{n,k}$: Under the promise that $\text{rank}(z) \in \{k, k+1\}$, compute $\text{rank}(z)$.
- Problem $\textsc{Inverse}_n$: Under the promise that z is invertible, decide whether the $(1,1)$ entry of z^{-1} is zero.
- Problem $\textsc{LinSolve}_{n,b}$: Under the promise that z is invertible, for a fixed non-zero vector $b \in \mathbb{F}_p^n$, consider the linear system $zt = b$ in the unknowns $t \in \mathbb{F}_p^n$. Decide whether t_1 is zero.

There are two natural ways to split z between Alice and Bob. In the concatenation model, Alice and Bob hold the top $n/2$ rows and bottom $n/2$ rows of z, respectively. In the additive split model, Alice and Bob hold $x, y \in M_n(\mathbb{F}_p)$ respectively, and $z = x + y$. The two models are equivalent up to a constant factor [31], see Section 5.3. All of this generalizes in the obvious manner to the multi-player setting.

Theorem 1. *Let f be one of* $\textsc{Rank}_{n,n-1}$, $\textsc{Inverse}_n$, *or* $\textsc{LinSolve}_{n,b}$. *Then* $\mathrm{R}_{1/10}(f) = \Omega(n^2 \log p)$.

The above immediately implies $\Omega(n^2 \log p)$ space lower bounds for randomized streaming algorithms for each of these problems, where the input matrix z is presented in row-major order. See Appendix C for details. Clearly these lower bounds are optimal, since the problems have trivial $O(n^2 \log p)$ upper bounds, that being the size of the input. We remark that Theorem 1 in fact extends to the *quantum* communication model, a generalization of randomized communication that we shall not elaborate on in this paper.

To prove these lower bounds, we use the *Fourier witness method* [31] for the promised rank problem, then reduce it to other problems. The reduction to the other problems critically uses the promise in the rank problem, for which establishing a lower bound was posed as an open question in [31]. Roughly speaking, the Fourier witness method is a special type of dual norm method [18, 28, 30]. In the dual norm method, there is a *witness* (a feasible solution of the dual maximization problem for the approximate norms). A typical choice of witness is the function itself (such as in the discrepancy method). In the *Fourier witness method* the witness is chosen as the Fourier transform of the function. This method works well for plus composed functions. For details, see Section 5.1.

We also consider the inner product and Hamming weight problems. Alice and Bob now hold *vectors* x and y.

- Problem IP_n: Under the promise that $\langle x, y \rangle \in \{0, 1\}$, compute $\langle x, y \rangle$. Here $x, y \in \mathbb{F}_p^n$.
- Problem $\textsc{Ham}_{n,k}$: Under the promise that $\|x + y\| \in \{k, k+2\}$, compute $\|x + y\|$. Here $x, y \in \mathbb{F}_2^n$ and $\|z\|$ denotes the Hamming weight of z, i.e., the number of 1 entries in z. Note that $x - y = x + y$.

We do not provide new two-player lower bounds for IP_n and $\textsc{Ham}_{n,k}$, but state the known ones here for use in our s-player lower bounds. It is known that $R_{1/3}(\mathrm{IP}_n) = \Omega(n \log p)$ [32], and $\mathrm{R}_{1/3}(\textsc{Ham}_{n,k}) = \Omega(k)$ [12].

s-Player Lower Bounds: For each of the above problems, there are natural s-player variants. We consider the *coordinator model* in which there is an additional player, called the coordinator, who has no input. We require that the s players can only talk to the coordinator. The message-passing model can be simulated in the coordinator model since every time a Player i wants to talk to a Player j, Player i can first send a message to the coordinator, and then the coordinator can forward the message to Player j. This only affects the communication by a factor of 2. See, e.g., Section 3 of [5] for a more detailed description.

For the matrix problems, Player i holds a matrix $x^{(i)}$ and the computations need to be performed on $z = x^{(1)} + \cdots + x^{(s)}$. The Hamming weight problem is similar, except that each $x^{(i)}$ is a vector in $\mathbb{F}_2^n$. For the inner product problem, each $x^{(i)} \in \mathbb{F}_p^n$ and we consider the generalized inner product, defined as $\sum_{j=1}^{n} \prod_{i=1}^{s} x_j^{(i)}$.

We provide a framework for applying the recent *symmetrization* technique of Phillips *et al.* [25] to each of these problems. Doing so lets us "scale up" each of the above lower bounds to the s-player versions of the problems.

However, the symmetrization technique of Phillips *et al.* does not immediately apply, since it requires a lower bound on the *distributional* communication complexity of the two-player problem under an input distribution with certain symmetric properties. Nevertheless, for many of the two-player lower bounds above, e.g., those in Theorem 1, our lower bound technique does not give a distributional complexity lower bound. We instead exploit the symmetry of the underlying problems, together with a re-randomization argument in Theorem 5 to argue that the hardest input distribution to these problems is in fact a symmetric distribution; see Definition 4 for a precise definition of symmetric. We thus obtain a distributional lower bound by the strong version of Yao's minimax principle.

We obtain the following results. Here, $\mathrm{R}^s_\delta(f)$ denotes the δ-error randomized communication complexity of the s-player variant of f. We give precise definitions in Section 3.

Theorem 2. *If f is one of* $\textsc{Rank}_{n,n-1}$, $\textsc{Inverse}_n$, *or* $\textsc{LinSolve}_{n,b}$, *then* $\mathrm{R}^s_{1/40}(f) = \Omega(sn^2 \log p)$. *Further,* $\mathrm{R}^s_{1/12}(\mathrm{IP}_n) = \Omega(sn \log p)$ *and* $\mathrm{R}^s_{1/12}(\textsc{Ham}_{n,k}) = \Omega(sk)$.

We note an application to the information-theoretic privacy of the $\textsc{rank}_{n,n-1}$ problem in Appendix D.

Related Work: Many linear algebra problems have been studied in both the communication complexity model and the streaming model. Chu and Schnitger [6] proved that $\Omega(n^2 \log p)$ communication is required by deterministic protocols for the singularity problem over $\mathbb{F}_p$. Luo and Tsitsiklis [19] proved that a deterministic protocol must transfer $\Omega(n^2)$ real numbers for the matrix inversion problem over $\mathbb{C}$, but Alice and Bob can only use addition, subtraction, multiplication and division of real numbers. Clarkson and Woodruff [7] proposed a randomized one pass streaming algorithm that uses $O(k^2 \log n)$ space to decide if the rank of an integer matrix is k and proved an $\Omega(k^2)$ lower bound for randomized one-way protocols in the communication complexity model via a reduction from the $\textsc{Indexing}$ communication problem. It implies an $\Omega(n^2)$ space lower bound in the streaming model with one pass. Miltersen et al. [20] showed a tight lower bound for deciding whether a vector is in a subspace of $\mathbb{F}_2^n$ in the one-sided error randomized asymmetric communication complexity model, using the Richness Lemma. Sun and Wang [31] proved the quantum communication complexities for matrix singularity and determinant over $\mathbb{F}_p$ are both $\Omega(n^2 \log p)$.

Compared to previous results, our results are stronger. For the rank problem, the matrix singularity problem in [31] is to decide if the rank of a matrix is n or less than n, but $\textsc{Rank}_{n,n-1}$ is to decide if the rank is n or $n-1$. This additional promise enables our lower bounds for $\textsc{Inverse}_n$ and $\textsc{LinSolve}_{n,b}$. If we set $k = n$ in Clarkson and Woodruff's result [7], the result gives us an $\Omega(n^2)$ bound for randomized one-way protocols. However, our lower bounds work even for quantum two-way protocols. For the inverse problem, Luo and Tsitsiklis's result [19] is in a non-standard communication complexity model, in which Alice and Bob can only make arithmetic operations on real numbers. However, our lower bound works in the standard model of communication complexity. A result of Miltersen et al. [20] is to decide if a vector is in a subspace. Sun and Wang [31] studied the problem deciding whether two $n/2$ dimensional subspaces intersect trivially (at $\{\mathbf{0}\}$ only) or not, but we get the same bound in Corollary 3 even with the promise. The results are analogous to the difference between set disjointness [3] and unique set disjointness [14, 26].

Corollary 3. *Alice and Bob each hold an $n/2$-dimensional subspace of $\mathbb{F}_p^n$. We promise that the intersection of the two subspaces is either $\{\mathbf{0}\}$ or a one-dimensional space. Any quantum protocol requires $\Omega(n^2 \log p)$*

communication to distinguish the two cases.

In the communication model, there is another way to distribute the input: Alice and Bob each hold an $n \times n$ matrix x and y, respectively, and they want to compute some property of $x+y$. This is equivalent to our model of matrix concatenation up to a constant factor, a fact we shall use in the paper (see [31] for a proof).

Paper Organization: In Section 3 we present our framework of multi-party communication lower bound for a class of problems. In Section 4 we discuss the IP_n problem and in Section 5 the $\textsc{Rank}_{n,n-1}$ problem and related linear algebra problems. Missing proofs, and the streaming and privacy applications are included in the Appendix.

2 Preliminaries

Communication Complexity: We briefly summarize the notions from communication complexity that we will need. For more background on communication complexity, we refer the reader to [16].

Let $f : X \times Y \to \{1,-1\}$ be a given function, which could be a partial function. Let $\mathrm{dom}(f)$ be the domain of definition of f. Alice and Bob, with unlimited computing power, want to compute $f(x,y)$ for $(x,y) \in \mathrm{dom}(f)$. Alice only knows $x \in X$ and Bob $y \in Y$. To perform the computation, they follow a protocol Π and send messages to each other in order to converge on a shared output $\Pi(x,y)$. We say a deterministic protocol Π computes f if $\Pi(x,y) = f(x,y)$ for all inputs $(x,y) \in \mathrm{dom}(f)$, and define the *deterministic communication complexity*, denoted by $D(f)$, to be the minimum over correct deterministic protocols for f, of the maximum number of bits communicated over all inputs. In a randomized protocol, Alice and Bob toss private coins and the messages can depend on the coin flips. We say a randomized protocol Π computes f with error probability δ if $\Pr\{\Pi(x,y) = f(x,y)\} \geq 1-\delta$ for all inputs $(x,y) \in \mathrm{dom}(f)$, and define the *randomized communication complexity*, denoted by $R_\delta(f)$, in the same way. When Alice and Bob share public random coins, the randomized communication complexity is denoted by $\mathrm{R}_\delta^{\mathrm{pub}}(f)$. Let μ be a probability distribution on $X \times Y$. The *μ-distributional communication complexity* of f, denoted by $D_\delta^\mu(f)$, is the least cost of a deterministic protocol for f with error probability at most δ with respect to μ. Yao's principle states that $R_\delta^{\mathrm{pub}}(f) = \max_\mu D_\delta^\mu(f)$.

In the model for multiparty communication complexity, there are s players, each gets an input $x_i \in X_i$, and they want to compute some function $f : X_1 \times \cdots \times X_s \to \{-1,1\}$ (which could be partially defined). We shall assume the coordinator model, in which there is an additional player called coordinator, who has no input. Players can only communicate with the coordinator but not each other directly. The coordinator will output the value of f. The private-coin, public-coin randomized communication complexity and μ-distributional communication complexity are denoted by $R_\delta^s(f)$, $R_\delta^{s,\mathrm{pub}}(f)$, and $D_\delta^{s,\mu}(f)$, respectively.

Information Theory: Let (X,Y) be a pair of discrete random variables with joint distribution $p(x,y)$. Suppose that X is a discrete random variable on Ω with distribution $p(x)$. Then the entropy $H(X)$ of the random variable X is defined by $H(X) = -\sum_{x\in\Omega} p(x)\log_2 p(x)$. The joint entropy $H(X,Y)$ of a pair of discrete random variables (X,Y) with joint distribution $p(x,y)$ is defined as $H(X,Y) = -\sum_x \sum_y p(x,y)\log p(x,y)$. The conditional entropy $H(X|Y)$ is defined as $H(X|Y) = \sum_y H(X|Y=y)\Pr\{Y=y\}$, where $H(X|Y=y)$ is the entropy of the conditional distribution of X given the event $\{Y=y\}$. The mutual information $I(X;Y)$ is defined as $I(X;Y) = \sum_{x,y} p(x,y)\log\frac{p(x,y)}{p(x)p(y)}$, where $p(x)$ and $p(y)$ are marginal distributions.

Information Cost: The following two definitions are from [10]. The *information cost* $\mathrm{ICost}(\Pi)$ of a protocol Π on input distribution μ equals the mutual information $I(X;\Pi(X))$, where X is a random variable distributed according to μ and $\Pi(X)$ is the transcript of Π on input X. The information complexity $\mathrm{IC}_{\mu,\delta}(f)$ of a problem f on a distribution μ with error probability δ is the infimum of $\mathrm{ICost}(\Pi)$ taken over all private-randomness protocols Π that err with probability at most δ for any input. When δ is clear from the context, we also write the information complexity as $\mathrm{IC}_\mu(f)$ for simplicity.

3 Reduction for Multi-Player Communication

Let $(G, \otimes)$ be a finite group and f be a function on G (could be a partial function). Suppose that $G = \bigcup_i G_i$ is the coarsest partition of G such that f is a constant function (allowing the value to be undefined) over each G_i. For a subset $X \subseteq G$, let $\mathrm{pre}(X) := \{(g_1, g_2) \in G \times G : g_1 \otimes g_2 \in X\}$. Let $I(f) = \{i : G_i \subseteq \mathrm{dom}(f)\}$, where $\mathrm{dom}(f) \subseteq G$ is the set on which f is defined.

We say that a family $\mathcal{H}$ of functions $h : G \times G \to G \times G$ is a *uniformizing family* for function f if there exists a probability measure μ on $\mathcal{H}$ such that for any i and $(g_1, g_2) \in \mathrm{pre}(G_i)$, when $h \in \mathcal{H}$ is randomly chosen according to μ, the image $h(g_1, g_2)$ is uniformly distributed on $\mathrm{pre}(G_i)$.

Example 1 ($\textsc{Rank}_{n,n-1}$). $G = M_n(\mathbb{F})$, the group of all $n \times n$ matrices over $\mathbb{F}$, with $\otimes$ being the usual matrix addition. In fact G is a ring, with the usual matrix multiplication. Define

$$f(x) = \begin{cases} 1, & \mathrm{rank}(x) = n; \\ 0, & \mathrm{rank}(x) = n-1; \\ \text{undefined}, & \text{otherwise}, \end{cases} \qquad x \in G.$$

Then $I(f) = \{1, 2\}$ and $G_1 = \{x \in G : \mathrm{rank}(x) = n\}$ and $G_2 = \{x \in G : \mathrm{rank}(x) = n-1\}$. The uniformizing family is $\mathcal{H} = \{h_{a,b}\}_{a \in G_1, b \in G}$ endowed with uniform measure, where $h_{a,b}(g_1, g_2) = (a(g_1 - b), a(g_2 + b))$.

Example 2 ($\textsc{Ham}_{n,k}$). $G = \mathbb{F}_2^n$ with the usual vector addition. Define

$$f(x) = \begin{cases} 1, & w(x) = k; \\ 0, & w(x) = k+2; \\ \text{undefined}, & \text{otherwise}, \end{cases} \qquad x \in G.$$

Then $|I(f)| = 2$. Let S_n denote the symmetric group of degree n. The uniformizing family $\mathcal{H} = \{h_{\sigma,b}\}_{\sigma \in S_n, b \in G}$ endowed with uniform measure, where $h_{\sigma,b}(g_1, g_2) = (\sigma(g_1 - b), \sigma(g_2 + b))$.

By reduction from Disjointness problem, we know that $\mathrm{R}^{\mathrm{pub}}_{1/10}(\textsc{Ham}_{k,k+2}) = \Omega(k)$.

As an auxiliary problem to the IP problem, we define

- Problem IP'_n: Suppose that $p > 2$. Alice and Bob hold two vectors $x, y \in (\mathbb{F}_p^*)^n$ respectively. We promise that inner product $\langle x, y \rangle \in \{0, 1\}$. They want to output $\langle x, y \rangle$.

Removing 0 from the scalar domain gives us a group structure as below.

Example 3 (IP'_n). $G = (\mathbb{F}_p^*)^n$ associated with the multiplication $\otimes$ defined to be the pointwise product, i.e., $x \otimes y = (x_1 y_1, x_2 y_2, \ldots, x_n y_n)$. Let $f(x) = \mathbf{1}_{\{x_1 + x_2 + \cdots + x_n = 0\}}$.

The following problem was considered in [32].

- Problem $\textsc{Cycle}_n$: Let π and σ be permutations in symmetric group S_n. Alice holds π and Bob σ, and they want to return 1 if $\pi \circ \sigma$ is exactly 1-cycle and return 0 otherwise.

Example 4 ($\textsc{Cycle}_n$). $G = S_n$, the symmetric group of degree n, with the usual permutation composition. Define

$$f(x) = \begin{cases} 1, & x \text{ has exactly one cycle}; \\ 0, & \text{otherwise}, \end{cases} \qquad x \in G.$$

Then $|I(f)| = 2$. The uniformizing family is $\mathcal{H} = \{h_{\sigma,\tau}\}_{\sigma,\tau \in S_n}$ endowed with uniform measure, where $h_{\sigma,\tau}(g_1, g_2) = (\sigma^{-1} g_1 \tau^{-1}, \tau g_2 \sigma)$. Observe that $g \mapsto \sigma^{-1} g \sigma$ maps a cycle $(a_1, \ldots, a_k)$ of g to $(\sigma(a_1), \ldots, \sigma(a_k))$, it is easy to verify that $\mathcal{H}$ is a uniformizing family indeed. It has been shown in [32] that $\mathrm{R}^{\mathrm{pub}}_{1/3}(\textsc{Cycle}_n) = \Omega(n)$.

We analyze the randomized communication complexity of problems that have a uniformizing family.

Definition 4. A distribution ν on $G \times G$ is called *weakly sub-uniform* if

1. ν is supported on $\bigcup_{i \in I(f)} \mathrm{pre}(G_i)$
2. $\nu|_{\mathrm{pre}(G_i)}$ is uniform for all $i \in I(f)$

In addition, if $\nu(\mathrm{pre}(G_i)) = 1/|I(f)|$ for all $i \in I(f)$, we say ν is the *sub-uniform* distribution.

Theorem 5. *If there exists a uniformizing family for f and $\delta \cdot |I(f)| < 1$, then for the two-player game computing f it holds that*

$$\mathrm{R}^{pub}_{\delta|I(f)|}(f) \le D^{\nu}_{\delta}(f) \le C \log_{|I(f)|\delta} \delta \cdot R^{pub}_{|I(f)|\delta}(f)$$

where $C > 0$ is an absolute constant and ν the sub-uniform distribution on $G \times G$.

Proof. Suppose the input is $(g_1, g_2) \in G \times G$. Next we describe a public-coin protocol Π'. With the public randomness, Alice and Bob choose a random h from the uniformizing family. They then run the optimal protocol Π_ν for input distribution ν (i.e., $\mathrm{cost}(\Pi_\nu) = D^{\nu}_{\delta}(f)$) on input $h(g_1, g_2)$.

It is not difficult to see that the public-coin protocol Π' has error probability at most $\delta \cdot |I(f)|$. Therefore, $R^{\mathrm{pub}}_{\delta \cdot |I(f)|} \le \mathrm{cost}(\Pi') = \mathrm{cost}(\Pi_\nu) = D^{\nu}_{\delta}(f)$. On the other hand, by Yao's principle, $\mathrm{R}^{\mathrm{pub}}_{\delta}(f) \ge D^{\nu}_{\delta}(f)$. Note that $R^{\mathrm{pub}}_{\delta}(f) \le C \log_{|I(f)|\delta} \delta \cdot R^{\mathrm{pub}}_{|I(f)|\delta}(f)$ for some absolute constant C, the conclusion follows. □

Now consider the following multi-player problem in coordinator model: There are s players and a coordinator. Each player receives an input $x_i \in G$. The coordinator will output the value of $f(x_1 \otimes x_2 \otimes \cdots \otimes x_s)$ with probability $\ge 1 - \delta$. Denote by $C^{s,\mathrm{pub}}_{\delta}(f)$ the number of bits that must be exchanged by the best protocol. By the symmetrization technique from [25], we have the following lemma.

Lemma 6. *Suppose that there exists a uniformizing family for f. Let ν be an arbitrary weakly sub-uniform distribution on $G \times G$ and Π_ν be a public-coin protocol that computes f with error probability δ on input distribution ν. Then $\mathrm{R}^{s,pub}_{\delta}(f) \ge s\mathbb{E}[\mathrm{cost}(\Pi_\nu)]$.*

Proof. Let ν_s be the distribution over G^s such that ν_s is the uniform distribution over $\mathrm{pre}_s(G_i) := \{(x_1, \ldots, x_s) \in G^s : x_1 \otimes \cdots \otimes x_s \in G_i\}$ when restricted onto it and $\nu_s(\mathrm{pre}_s(G_i)) = \nu(\mathrm{pre}(G_i))$. Let Π_s be an s-player (deterministic) protocol for input distribution ν_s with error probability δ.

Consider the following two-player protocol Π' on input $(g_1, g_2) \sim \nu$: First suppose that Alice and Bob have public randomness. They first use the public randomness to agree on an index j chosen at random uniformly from $\{1, \ldots, s\}$. Alice also generates, using her own randomness, the input $\{x_i\}_{i \ne j}$ of other players uniformly at random conditioned on $\bigotimes_{i \ne j} x_i = g_1$. Then Alice and Bob run the s-player protocol, in which Bob simulates player j with input $x_j := g_2$, and Alice simulates all other players and the coordinator. The message sent in this protocol is just the message sent between the coordinator and player j in Π_s.

It is not hard to see that $(x_1, \ldots, x_s) \sim \nu_s$. It follows from a symmetrization argument like the proof [25, Theorem 1.1] that $\mathbb{E}[\mathrm{cost}(\Pi')] \le \mathrm{cost}(\Pi_s)/s$, where the expectation is taken over the public coins. The conclusion follows from taking the infimum over Π_s. □

Theorem 7. *Suppose that there exists a uniformizing family for f, then $\mathrm{R}^{s,pub}_{\delta}(f) \ge \delta s \, \mathrm{R}^{pub}_{2|I(f)|\delta}(f)$.*

Proof. Pick ν to be the sub-uniform distribution in the preceding lemma. By fixing the public coins and a Markov bound, one can construct a two-player deterministic protocol Π'' such that $\mathrm{cost}(\Pi'') \le (1/\delta)\,\mathrm{cost}(\Pi_s)/s$ and Π'' succeeds with probability at least $1 - 2\delta$ when the input is distributed as ν. Hence $D^{\nu}_{2\delta}(f) \le (1/\delta)\,\mathrm{cost}(\Pi_s)/s$. It then follows from Theorem 5 that $R^{\mathrm{pub}}_{2|I(f)|\delta}(f) \le (1/\delta)\,\mathrm{cost}(\Pi_s)/s$. Taking infimum over Π_s, we obtain that $\mathrm{R}^{\mathrm{pub}}_{2|I(f)|\delta}(f) \le (1/\delta) D^{\nu_s}_{\delta}(f)/s \le (1/\delta) R^{s,\mathrm{pub}}_{\delta}(f)/s$. □

The following are immediate corollaries of the theorem above applied to our previous Example 2 and 4. We leave the results of Example 1 and 3 for later sections.

Corollary 8. $\mathrm{R}^{s,pub}_{1/12}(\textsc{Ham}_{k,k+2}) = \Omega(s\,\mathrm{R}^{pub}_{1/3}(\textsc{Ham}_{k,k+2})) = \Omega(sk)$.

Corollary 9. $\mathrm{R}^{s,pub}_{1/12}(\textsc{Cycle}_n) = \Omega(sn)$.

4 The IP Problem

Let p be a prime. Sun *et al.* considered a variant of the IP problem, denoted by IP''_n, in which Alice has $x \in F_p^n$ and Bob $y \in (F_p^*)^n$, and showed that $\mathrm{R}^{\mathrm{pub}}_{1/3}(\mathrm{IP}''_n) = \Omega(n \log p)$ [32]. Via a simple reduction, we show that

Lemma 10. *When $p \geq p_0$ for some constant p_0, it holds that* $\mathrm{R}^{pub}_{1/3}(\mathrm{IP}'_n) = \Omega(n \log p)$.

Proof. For an input of IP'', Alice can send the indices of the zero coordinates to Bob using n bits; on the remaining coordinates, Alice and Bob have an instance of IP' of size at most n. Hence $n + \mathrm{R}^{\mathrm{pub}}_{1/3}(\mathrm{IP}'_n) \geq \mathrm{R}^{\mathrm{pub}}_{1/3}(\mathrm{IP}''_n)$, whence the conclusion follows. □

It is clear, by Yao's principle, that $\mathrm{R}^{\mathrm{pub}}_{\delta}(\mathrm{IP}_n) \geq \mathrm{R}^{\mathrm{pub}}_{\delta}(\mathrm{IP}'_n)$. Now, as an immediate corollary of Theorem 7, we have

Theorem 11. $\mathrm{R}^{s,pub}_{1/12}(\mathrm{IP}_n) = \Omega(sn \log p)$.

Proof. Let p_0 be as in Lemma 10. It follows from Lemma 10 and Theorem 7 that $\mathrm{R}^{\mathrm{s,pub}}_{1/12}(\mathrm{IP}) \geq \mathrm{R}^{\mathrm{s,pub}}_{1/12}(\mathrm{IP}') = \Omega(s\,\mathrm{R}^{\mathrm{pub}}_{1/3}(\mathrm{IP}')) = \Omega(sn \log p)$. When $p < p_0$, the result is due to Braverman et al. in [5], who prove an $\Omega(sn)$ lower bound for IP over the integers with the promise that the inner product is 0 or 1. Note that this implies an $\Omega(sn \log p)$ lower bound for computing IP over $\mathbb{F}_p$ as well, since $p < p_0$ is a fixed constant. □

5 The Rank Problem

We shall use the Fourier witness method to prove a lower bound on $\textsc{Rank}'_{n,n-1}$. We then use this result for $\textsc{Rank}_{n,n-1}$ to obtain lower bounds for the other problems. We review some basics of the Fourier witness method in Section 5.1 then give the proof of the lower bound in Section 5.2.

5.1 Fourier Witness Method

5.1.1 Fourier Analysis

For prime p, let $\mathbb{F}_p$ be the finite field of p elements. We define the Fourier transformation on the group $(\mathbb{F}_p^n, +)$.

Definition 12 (Fourier transform)**.** Let $f : \mathbb{F}_p^N \to \mathbb{R}$ be a function. Then, the Fourier coefficient of f, denoted by $\hat{f}$, is also a $\mathbb{F}_p^N \to \mathbb{R}$ function, defined as $\hat{f}(s) = \frac{1}{p^N} \sum_{x \in \mathbb{F}_p^N} \omega^{-\langle s,x \rangle} f(x)$, where $\omega = e^{2\pi i/p}$.

Fact 13. $f = p^N \widehat{\left((\hat{f})^*\right)}^*$.

5.1.2 Approximate Norm and Dual Norm

The ℓ_p norm of a vector $v \in \mathbb{R}^n$ is defined by $\|v\|_p := \left(\sum_{i=1}^n |v_i|^p\right)^{1/p}$ and the ℓ_∞ norm by $\|v\|_\infty := \max_{i=1}^n |v_i|$. The trace norm of an $n \times n$ matrix F, denoted by $\|F\|_{\mathrm{tr}}$, is defined as $\|F\|_{\mathrm{tr}} := \sum_i \sigma_i$, where $\sigma_1, \cdots, \sigma_n$ are the singular values of F.

The matrix rank and some matrix norms can give lower bounds for deterministic communication complexity. For randomized lower bounds, we need the notions of approximate rank and norms.

Definition 14 (approximate norm)**.** Let $\rho : \mathbb{R}^X \mapsto \mathbb{R}$ be an arbitrary norm and $f : X \mapsto \mathbb{R}$ a partial sign function. The ε-approximate ρ norm of f, denoted by $\rho^\varepsilon(f)$, is defined as $\rho^\varepsilon(f) = \inf_\phi \rho(\phi)$, where the infimum is taken over all functions $\phi : X \mapsto \mathbb{R}$ that satisfy

$$\phi(x) \in \begin{cases} [1-\varepsilon, 1+\varepsilon] & \text{if } f(x) = 1; \\ [-1-\varepsilon, -1+\varepsilon] & \text{if } f(x) = -1; \\ [-1-\varepsilon, 1+\varepsilon] & \text{if } f(x) \text{ is undefined.} \end{cases}$$

The following lemma shows that the approximate trace norm gives lower bounds on quantum communication complexity, as well as on randomized protocols with public coins. The following lemma is a result in [17] combined with Neumann's argument for converting a public-coin protocol into a private-coin one.

Lemma 15. *For $\delta > 0$ such that $1/(1-2\varepsilon) \le 1+\delta$, it holds that*

$$\mathrm{R}_\delta^{pub}(f) \ge \Omega\left(\log \frac{(\|F\|_{\mathrm{tr}}^\varepsilon)^2}{\mathrm{size}(F)}\right) - O(\log n + \log \frac{1}{\delta}).$$

The approximate norms are minimization problems. We will consider the dual problems, which are maximization problems.

Definition 16. Let ρ be an arbitrary norm on $\mathbb{R}^n$. The dual norm of ρ, denoted by ρ^*, is defined as

$$\rho^*(v) = \sup_{u:\rho(u)\le 1} \langle v, u \rangle.$$

The following lemma characterizes the approximate norm as a maximization problem so that we can prove *lower* bounds more easily.

Lemma 17 ([29])**.** *Let f be a partial sign function and ρ an arbitrary norm. Then*

$$\rho^\varepsilon(f) = \sup_{\psi \ne 0} \frac{\langle f, \psi \cdot \mathrm{dom}(f)\rangle - \|\psi \cdot \overline{\mathrm{dom}(f)}\|_1 - \varepsilon\|\psi\|_1}{\rho^*(\psi)}, \quad \varepsilon > 0.$$

where

$$\mathrm{dom}(f)(x) = \begin{cases} 1 & \textit{if } f(x) \textit{ is defined,} \\ 0 & \textit{otherwise,} \end{cases}$$

$\overline{\mathrm{dom}(f)}(x) = 1 - \mathrm{dom}(f)(x)$*, and* $(\psi \cdot \varphi)(x) = \psi(x)\varphi(x)$*.*

We call a feasible solution in the dual problem the witness of the original problem. In particular, in Lemma 17, the function ψ is the witness. Any ψ gives a lower bound for $\rho^\varepsilon(f)$. It is difficult to find a useful witness. The first choice that comes to mind is to choose $\psi = f \cdot \mathrm{dom}(f)$, because it makes $\langle f, \psi \cdot \mathrm{dom}(f)\rangle$ large and $\|\overline{\mathrm{dom}(f)}\|_1$ small. This is the discrepancy method. We use a different choice: $\psi = \widehat{(f \cdot \overline{\mathrm{dom}(f)})}$. We call it the *Fourier witness method*, introduced in [31], but used here for partial functions.

Definition 18 (approximate Fourier p-norm)**.** Let $f : \mathbb{F}_p^N \mapsto \mathbb{R}$ be a function and $p \geq 1$. The Fourier p-norm of f, denoted by $\|\hat{f}\|_p$, is the p-norm of $\hat{f}$. Furthermore, if f is a sign function, the approximate Fourier p-norm of f, denoted by $\|\hat{f}\|_p^\varepsilon$, is the approximate $\|\hat{\cdot}\|_p$ norm of f.

Fact 19. *The dual norm of* $\|\cdot\|_1$ *is* $\|\cdot\|_\infty$*. The dual norm of* $\|\hat{\cdot}\|_1$ *is* $p^N\|\hat{\cdot}\|_\infty$*.*

The Fourier coefficients of a plus composed function are related to the singular values of the associated matrix, as shown by the following lemma, whose proof is postponed to Appendix A. Applied to approximate trace norm, it also builds a bridge between the approximate trace norm and the approximate Fourier ℓ_1-norm for a plus composed function. Similar results and additional background can be found in [17, 31].

Lemma 20. *Suppose that* $g : \mathbb{F}_p^N \mapsto \mathbb{R}$ *is a function, and* f *is a plus-composed function* $f(x, y) = g(x + y)$*. Let* F *be the associated matrix of* f*. Then the singular values of* F *are* p^N *times the modulus of the Fourier coefficients of* g*, i.e.* $\sigma_F = p^N \cdot |\hat{g}|$*, where* σ_F *are the singular values of* F *and* $|\hat{g}|(s) = |\hat{g}(s)|$*. As a consequence,* $\|F\|_{\mathrm{tr}}^\varepsilon = p^N \cdot \|\hat{g}\|_1^\varepsilon$*.*

5.2 $\textsc{Rank}_{n,n-1}$

For a matrix $x \in \mathbb{F}_p^{n\times n}$, we define $\theta(x) = 1$ if x is of full rank and $\theta(x) = 0$ otherwise. We shall use θ as the witness in the proof of $\textsc{Rank}'_{n,n-1}$. The same function θ has been used to prove a communication complexity lower bound for the matrix singularity problem in [31].

Theorem 21. $\mathrm{R}_{1/10}^{pub}(\textsc{Rank}_{n,n-1}) = \Omega(n^2 \log p)$.

Proof. Suppose that Π is a public-coin protocol for $\textsc{Rank}'_{n,n-1}$ with error probability $\leq 1/10$. Then Alice and Bob can build a public-coin protocol Π' as follows. They use the public coins to choose a random matrix r and run Π on input $(x - r, y + r)$. It is easy to see that Π' has error probability $\leq 1/10$ and $\mathrm{cost}(\Pi') = \mathrm{cost}(\Pi)$. Observe that the distribution of $\Pi'(x, y)$ is identical to the that of $\Pi'(a, b)$ whenever $x + y = a + b$.

Define the partial sign function $g(x) = 1$ if $\mathrm{rank}(x) = n$, $g(x) = -1$ if $\mathrm{rank}(x) = n - 1$, and $g(x)$ is undefined otherwise. Let $f(x, y)$ be the expected output of $\Pi'(x, y)$. Then f is a plus-composed function. By the correctness of Π, we know that $f(x, y) = g(x + y)$ whenever $g(x + y)$ is defined. We claim that $\|g\|_1^\varepsilon = \Omega(p^{n(n-3)/2})$ for $\varepsilon = 1/4$, following Lemma 17 (applied with witness θ as in the paragraph before the theorem statement) and Fact 19. See Appendix B for details. Finally, it follows from Lemma 15 that

$$\mathrm{R}_{1/10}^{\mathrm{pub}}(f) = \Omega\left(\log \frac{\|F\|_{\mathrm{tr}}^{1/4}}{\sqrt{\mathrm{size}(F)}}\right) - O(\log n) = \Omega\left(\log \frac{p^{n^2}\|\hat{g}\|_1^\varepsilon}{\sqrt{\mathrm{size}(F)}}\right) - O(\log n)$$
$$= \Omega\left(\log \frac{p^{n^2} \cdot 0.4p^{n(n-3)/2}}{p^{n^2}}\right) - O(\log n) = \Omega(n^2 \log p). \quad \square$$

The lower bound for the multi-player $\textsc{Rank}$ problem is an immediate corollary of Theorem 7.

Corollary 22. $R_{1/40}^{s,pub}(\textsc{Rank}_{n,n-1}) = \Omega(sR_{1/10}^{pub}(\textsc{Rank}_{n,n-1})) = \Omega(sn^2 \log p)$.

By padding zeros outside the top-left $k \times k$ submatrix, we obtain a lower bound for $\textsc{Rank}_{k,k-1}$.

Corollary 23. $\mathrm{R}_{1/10}^{pub}(\textsc{Rank}_{k,k-1}) = \Omega(k^2 \log p)$.

5.3 Linear Algebra Problems

Problem 1 ($\textsc{Inverse}$)**.** *Alice and Bob hold two* $n \times n$ *matrices* x *and* y *over* $\mathbb{F}_p$*, respectively. We promise that* $x + y$ *is invertible over* $\mathbb{F}_p$*. They want to determine if the top-left entry of* $(x + y)^{-1}$ *is zero (output* -1*) or non-zero (output* 1*).*

Problem 2 (LinSolve)**.** *Alice and Bob hold two $n \times n$ matrices x and y over $\mathbb{F}_p$, respectively. We promise that $x + y$ is invertible over $\mathbb{F}_p$. b is a parameter of this problem. t is the vector of variables of the linear system $(x + y)t = b$. They want to determine if the first coordinate of t is zero.*

Theorem 24. $\mathrm{R}^{pub}_{1/20}(\textsc{Inverse}) = \Omega(n^2 \log p)$ *for* $p \geq 3$.

Proof. We reduce Rank to Inverse. Let $A = x + y$ and $\tilde{A}$ be the lower-right $(n-1) \times (n-1)$ block of A. Then $A^{-1}_{11} = 0$ iff $\mathrm{rank}(\tilde{A}) < n - 1$.

Now, suppose that A is an $(n-1) \times (n-1)$ matrix and $\mathrm{rank}(A) \in \{n-1, n-2\}$. We augment A to A_1 by appending a random column. With probability $1 - 1/p$ it holds that $\mathrm{rank}(A_1) = n - 1$ when $\mathrm{rank}(A) = n - 2$. Now we augment A_1 to A_2 by appending a random row. With probability $1 - 1/p$ it holds that $\mathrm{rank}(A_2) = n$ when $\mathrm{rank}(A_1) = n - 1$.

Run a protocol for Inverse on A_2. We denote the communication complexity of the protocol by $c(n)$. When $\mathrm{rank}(A) = n - 1$, if the error probability of the protocol is at most $1/20$, then it outputs 1 with probability $\alpha \leq \frac{1}{20}(1 - \frac{1}{p}) + \frac{1}{p}$, while when $\mathrm{rank}(A) = n - 2$ it outputs 1 with probability $\beta \geq \frac{19}{20}(1 - \frac{1}{p})^2$. Then $\beta - \alpha \geq \frac{19}{20}(1 - \frac{1}{p})^2 - \frac{1}{20}(1 - \frac{1}{p}) - \frac{1}{p} \geq \frac{1}{18}, \quad p \geq 3$, which implies that $\Theta(1)$ independent repetitions allow us to solve Rank on $(n-1) \times (n-1)$ matrices, i.e., to distinguish $\mathrm{rank}(A) = n - 1$ from $\mathrm{rank}(A) = n - 2$, with error probability $\leq 1/20$ and communication complexity $\Theta(c(n)) = \Omega((n-1)^2 \log p)$. Therefore $c(n) = \Omega((n-1)^2 \log p) = \Omega(n^2 \log p)$. □

Theorem 25. $\mathrm{R}^{pub}_{1/20}(\textsc{Inverse}) = \Omega(n^2)$ *for* $p = 2$.

Proof. As before, we augment A to A_2. Here we further randomize A_2 by multiplying a random invertible matrix on both sides of A_2, that is, we form $B = G_1 A_2 G_2$ where G_1, G_2 are uniform over $n \times n$ non-singular matrices over $\mathbb{F}_p$. It is clear that $\mathrm{rank}(B) = \mathrm{rank}(A_2)$, and B is uniformly distributed over the $n \times n$ matrices with the same rank.

Run a protocol for Inverse on B. Suppose that it outputs zero with probability p_0 when the input matrix has rank $n - 1$. This probability can be calculated by Alice and Bob individually with no communication cost. When $\mathrm{rank}(A) = n - 1$, it outputs 1 with probability $\alpha = \frac{1}{20}(1 - \frac{1}{p}) + \frac{p_0}{p} = \frac{1}{40} + \frac{p_0}{2}$, while when $\mathrm{rank}(A) = n - 2$ it outputs 1 with probability $\beta \geq \frac{19}{20}(1 - \frac{1}{q})^2 + p_0 \cdot \frac{2}{p}(1 - \frac{1}{p}) = \frac{19}{80} + \frac{p_0}{2}$. Then, $\alpha - \beta \geq \frac{17}{80}$. The rest follows as in the proof for $p \geq 3$. □

Now we reduce Inverse to $\textsc{LinSolve}_b$.

Theorem 26. $\mathrm{R}^{pub}_{1/20}(\textsc{LinSolve}_b) = \Omega(n^2 \log p)$ *for* $b \neq \mathbf{0}$,

Proof. We prove it by a reduction from Inverse.

Take an instance (x, y) from Inverse. Since $b \neq 0$, there exists an invertible matrix Q such that $Qb = (1, 0, 0, \cdots, 0)^{\mathrm{T}}$. Alice and Bob agree on the same Q, e.g. the minimal Q in alphabetical order. Then they run the protocol of LinSolve on input $(Q^{-1}x, Q^{-1}y, b)$. Then $t = (Q^{-1}x + Q^{-1}y)^{-1}b = (x + y)^{-1}QQ^{-1}(1, 0, \cdots, 0)^{\mathrm{T}} = (x + y)^{-1}(1, 0, \cdots, 0)^{\mathrm{T}}$ and thus $t_1 = ((x + y)^{-1})_{11}$. □

Acknowledgements We would like to thank Amit Chakrabarti for reading and helping with an earlier draft of this paper, and Troy Lee for giving us useful suggestions. David Woodruff would also like to acknowledge the XDATA program of the Defense Advanced Research Projects Agency (DARPA), administered through Air Force Research Laboratory contract FA8750-12-C0323, for support for this project.

References

[1] Anil Ada, Arkadev Chattopadhyay, Stephen Cook, Lila Fontes, Michal Koucky, and Toniann Pitassi. The hardness of being private. In *Proceedings of the 2012 IEEE Conference on Computational Complexity (CCC)*, CCC '12, pages 192–202, 2012.

[2] Noga Alon, Yossi Matias, and Mario Szegedy. The Space Complexity of Approximating the Frequency Moments. *J. Comput. Syst. Sci.*, 58(1):137–147, 1999.

[3] Laszlo Babai, Peter Frankl, and Janos Simon. Complexity classes in communication complexity theory. In *Foundations of Computer Science, 1986., 27th Annual Symposium on*, pages 337–347, 1986.

[4] Boaz Barak, Mark Braverman, Xi Chen, and Anup Rao. How to compress interactive communication. In *Proceedings of the 42Nd ACM Symposium on Theory of Computing*, STOC '10, pages 67–76, 2010.

[5] Mark Braverman, Faith Ellen, Rotem Oshman, Toniann Pitassi, and Vinod Vaikuntanathan. A tight bound for set disjointness in the message-passing model. In *FOCS*, pages 668–677, 2013.

[6] J. I. Chu and G. Schnitger. Communication complexity of matrix computation over finite fields. *Theory of Computing Systems*, 28:215–228, 1995. 10.1007/BF01303056.

[7] Kenneth L. Clarkson and David P. Woodruff. Numerical linear algebra in the streaming model. In *Proceedings of the 41st annual ACM symposium on Theory of computing*, STOC '09, pages 205–214, New York, NY, USA, 2009. ACM.

[8] Graham Cormode, S. Muthukrishnan, and Ke Yi. Algorithms for distributed functional monitoring. In *Proc. 19th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1076–1085, 2008.

[9] Alan Frieze, Ravi Kannan, and Santosh Vempala. Fast Monte-Carlo algorithms for finding low-rank approximations. In *Foundations of Computer Science, 1998. Proceedings. 39th Annual Symposium on*, pages 370–378. IEEE, 1998.

[10] Venkatesan Guruswami and Krzysztof Onak. Superlinear lower bounds for multipass graph processing. In *IEEE Conference on Computational Complexity*, pages 287–298, 2013.

[11] Monika Rauch Henzinger, Prabhakar Raghavan, and Sridar Rajagopalan. Computing on data streams. In *External Memory Algorithms: Dimacs Workshop External Memory and Visualization, May 20-22, 1998*, volume 50, page 107. American Mathematical Soc., 1999.

[12] Wei Huang, Yaoyun Shi, Shengyu Zhang, and Yufan Zhu. The communication complexity of the hamming distance problem. *Inf. Process. Lett.*, 99(4):149–153, 2006.

[13] Rahul Jain and Hartmut Klauck. The partition bound for classical communication complexity and query complexity. In *IEEE Conference on Computational Complexity*, pages 247–258, 2010.

[14] Bala Kalyanasundaram and Georg Schnitger. The Probabilistic Communication Complexity of Set Intersection. *SIAM J. Discrete Math.*, 5(4):545–557, 1992.

[15] Iordanis Kerenidis, Sophie Laplante, Virginie Lerays, Jeremie Roland, and David Xiao. Lower bounds on information complexity via zero-communication protocols and applications. In *Proceedings of the 2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*, FOCS '12, pages 500–509, 2012.

[16] E. Kushilevitz and N. Nisan. *Communication Complexity*. Cambridge Univ Pr, 1997.

[17] Troy Lee and Adi Shraibman. Lower Bounds in Communication Complexity. *Foundations and Trends in Theoretical Computer Science*, 3(4):363–399, October 2009.

[18] Nati Linial and Adi Shraibman. Lower bounds in communication complexity based on factorization norms. *Random Structures & Algorithms*, 34(3):368–394, 2009.

[19] Zhi-Quan Luo and John N. Tsitsiklis. On the communication complexity of distributed algebraic computation. *J. ACM*, 40:1019–1047, November 1993.

[20] Peter Bro Miltersen, Noam Nisan, Shmuel Safra, and Avi Wigderson. On Data Structures and Asymmetric Communication Complexity. *J. Comput. Syst. Sci.*, 57(1):37–49, 1998.

[21] K Morrison. Integer sequences and matrices over finite fields. *J. Integer Seq*, 9(2):06–2, 2006.

[22] J. I. Munro and M. S. Paterson. Selection and sorting with limited storage. *Theoretical Computer Science*, 12(3):315 – 323, 1980.

[23] S Muthukrishnan. *Data streams: Algorithms and applications*. Now Publishers Inc, 2005.

[24] Ilan Newman. Private vs. common random bits in communication complexity. *Information Processing Letters*, 39(2):67–71, 1991.

[25] Jeff M. Phillips, Elad Verbin, and Qin Zhang. Lower bounds for number-in-hand multiparty communication complexity, made easy. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '12, pages 486–501, 2012.

[26] Alexander A. Razborov. On the distributional complexity of disjointness. *Theoretical Computer Science*, 106(2):385–390, 1992.

[27] Tamas Sarlos. Improved approximation algorithms for large matrices via random projections. In *Foundations of Computer Science, 2006. 47th Annual IEEE Symposium on*, pages 143–152. IEEE, 2006.

[28] Alexander A. Sherstov. The pattern matrix method for lower bounds on quantum communication. In *Proceedings of the 40th annual ACM symposium on Theory of computing*, STOC '08, pages 85–94, New York, NY, USA, 2008. ACM.

[29] Alexander A Sherstov. Strong direct product theorems for quantum communication and query complexity. *SIAM Journal on Computing*, 41(5):1122–1165, 2012.

[30] Yaoyun Shi and Yufan Zhu. Quantum communication complexity of block-composed functions. *Quantum Information and Computation*, 9:444–460, May 2009.

[31] Xiaoming Sun and Chengu Wang. Randomized Communication Complexity for Linear Algebra Problems over Finite Fields. In Christoph Dürr and Thomas Wilke, editors, *29th International Symposium on Theoretical Aspects of Computer Science (STACS 2012)*, Leibniz International Proceedings in Informatics (LIPIcs), 2012.

[32] Xiaoming Sun, Chengu Wang, and Wei Yu. The relationship between inner product and counting cycles. In *Proceedings of the 10th Latin American International Conference on Theoretical Informatics*, LATIN'12, pages 643–654, 2012.

[33] David P. Woodruff and Qin Zhang. Tight bounds for distributed functional monitoring. In *STOC*, pages 941–960, 2012.

[34] David P. Woodruff and Qin Zhang. When distributed computation is communication expensive. In *DISC*, pages 16–30, 2013.

[35] David P. Woodruff and Qin Zhang. An optimal lower bound for distinct elements in the message passing model. In *SODA*, pages 718–733, 2014.

[36] Andrew Chi-Chih Yao. Some complexity questions related to distributive computing (Preliminary Report). In *Proceedings of the eleventh annual ACM symposium on Theory of computing*, pages 209–213. ACM, 1979.

A Proof of Lemma 20

Proof. Let $\omega = e^{2\pi/p}$, $U(x,y) = p^{-N/2} \cdot \omega^{-\langle x,y\rangle}$, and $\Lambda = \mathrm{diag}(|\hat{g}|^2)$. We will prove $F^\dagger F = p^N U \Lambda U^\dagger$.

$$\begin{aligned}
(F^\dagger F)_{x,z} &= \sum_y (F_{y,x})^* F_{y,z} \\
&= \sum_y g(y+x)^* g(y+z) \\
&= \sum_y \sum_s \omega^{-\langle s, y+x\rangle} \hat{g}(s)^* \sum_t \omega^{\langle t, y+z\rangle} \hat{g}(t) \\
&= \sum_s \sum_t \omega^{-\langle s,x\rangle + \langle t,z\rangle} \hat{g}(s)^* \hat{g}(t) \left(\sum_y \omega^{\langle -s+t, y\rangle} \right) \\
&= \sum_s \sum_t \omega^{-\langle s,x\rangle + \langle t,z\rangle} \hat{g}(s)^* \hat{g}(t) \cdot p^N \delta_{s,t} \\
&= p^N \cdot \sum_s \omega^{-\langle s,x\rangle + \langle s,z\rangle} \hat{g}(s)^* \hat{g}(s) \\
&= p^N \cdot \sum_s \omega^{-\langle s,x\rangle} |\hat{g}(s)|^2 \omega^{\langle s,z\rangle} \\
&= p^{2N} \cdot \sum_s U_{x,s} |\hat{g}(s)|^2 (U_{z,s})^* \\
&= p^{2N} (U \Lambda U^\dagger)_{x,z}
\end{aligned}$$

U is unitary, because

$$\begin{aligned}
(U^\dagger U)_{x,z} &= \sum_y (U_{y,x})^* U_{y,z} \\
&= \sum_y p^{-N/2} \omega^{\langle y,x\rangle} \cdot p^{-N/2} \omega^{-\langle y,z\rangle} \\
&= p^{-N} \cdot \sum_y \omega^{\langle y, x-z\rangle} \\
&= p^{-N} \cdot p^N \delta_{x,z} \\
&= \delta_{x,z}
\end{aligned}$$

Therefore, the singular values of F are $p^N \cdot |\hat{g}|$. □

B Proofs for the RANK$_{n,n-1}$ Problem

We shall use the following fact.

Fact 27 ([21])**.** *The number $n \times n$ matrices over $\mathbb{F}_p$ of rank-r is*

$$\frac{(p^n-1)(p^n-p)\cdots(p^n-p^{r-1})}{(p^r-1)(p^r-p)\cdots(p^r-p^{r-1})} \prod_{k=0}^{r-1} (p^n - p^k).$$

The following lemma computes the Fourier coefficients and shows that the ℓ_1-norm of $\hat{\theta}$ is small.

Lemma 28. *Let* $r = \mathrm{rank}(s)$*, then*

$$\hat{\theta}(s) = (-1)^r p^{-n(n+1)/2} \prod_{k=1}^{n-r} (p^k - 1).$$

Hence

$$\|\hat{\theta}\|_1 = p^{-n} \prod_{k=1}^{n} (p^k - 1) \prod_{k=0}^{n-1} \frac{1 + p^k}{p^k}.$$

Proof. The expression of $\hat{\theta}(s)$ is from [31]. It follows straightforwardly that

$$\begin{aligned}
\|\hat{\theta}\|_1 &= \sum_{r=0}^{n} \sum_{s:\mathrm{rank}(s)=r} |\hat{\theta}(s)| \\
&= \sum_{r=0}^{n} p^{r(r-1)/2} \binom{n}{r}_p \prod_{k=n-r+1}^{n} (p^k - 1) \cdot p^{-n(n+1)/2} \prod_{k=1}^{n-r} (p^k - 1) \\
&= p^{-n(n+1)/2} \prod_{k=1}^{n} (p^k - 1) \sum_{r=0}^{n} p^{r(r-1)/2} \binom{n}{r}_p \\
&= p^{-n(n+1)/2} \prod_{k=1}^{n} (p^k - 1) \prod_{k=0}^{n-1} (1 + p^k) \\
&= p^{-n} \prod_{k=1}^{n} (p^k - 1) \prod_{k=0}^{n-1} \frac{1 + p^k}{p^k}
\end{aligned}$$

□

We are now ready to complete the proof of Theorem 21.

Proof. We shall show that $\|\hat{g}\|_1^{\varepsilon}$ is large. By Lemma 17 and Fact 19,

$$\|\hat{g}\|_1^{\varepsilon} \geq \frac{\langle g, \psi \cdot \mathrm{dom}(g) \rangle - \|\psi \cdot \overline{\mathrm{dom}(g)}\|_1 - \varepsilon \|\psi\|_1}{p^{n^2} \|\hat{\psi}\|_\infty}.$$

Choosing $\psi = (-1)^n \hat{\theta}$ yields that

$$\begin{aligned}
&\langle g, \psi \cdot \mathrm{dom}(g) \rangle \\
&= \sum_{x:\mathrm{rank}(x)=n} (-1)^n \hat{\theta}(x) - \sum_{x:\mathrm{rank}(x)=n-1} (-1)^n \hat{\theta}(x) \\
&= \sum_{x:\mathrm{rank}(x)=n} |\hat{\theta}(x)| + \sum_{x:\mathrm{rank}(x)=n-1} |\hat{\theta}(x)| \\
&= p^{n(n-1)/2} \prod_{k=1}^{n} (p^k - 1) \cdot p^{-n(n+1)/2} + \frac{p^n - 1}{(p-1)^2} p^{(n-1)(n-2)/2} \prod_{k=1}^{n} (p^k - 1) \cdot p^{-n(n+1)/2} (p - 1) \\
&= \left(1 + \frac{p - p^{-n+1}}{p - 1}\right) p^{-n} \prod_{k=1}^{n} (p^k - 1).
\end{aligned}$$

Observe that

$$\|\psi \cdot \mathrm{dom}(g)\|_1 = \sum_{x:\mathrm{rank}(x)=n} |\hat{\theta}(x)| + \sum_{x:\mathrm{rank}(x)=n-1} |\hat{\theta}(x)| = \langle g, \psi \cdot \mathrm{dom}(g)\rangle$$

$$\|\psi \cdot \overline{\mathrm{dom}(g)}\|_1 = \|\psi\|_1 - \|\psi \cdot \mathrm{dom}(g)\|_1 = \|\hat{\theta}\|_1 - \|\psi \cdot \mathrm{dom}(g)\|_1$$

$$\|\psi\|_1 = \|\hat{\theta}\|_1$$

$$\|\hat{\psi}\|_\infty = \|\hat{\hat{\theta}}\|_\infty = p^{-n^2}.$$

The lower bound for $\|\hat{g}\|_1^\varepsilon$ follows as below.

$$\begin{aligned}
\|\hat{g}\|_1^\varepsilon &\geq \frac{\langle g, \psi \cdot \mathrm{dom}(g)\rangle - \|\psi \cdot \overline{\mathrm{dom}(g)}\|_1 - \varepsilon\|\psi\|_1}{p^{n^2}\|\hat{\psi}\|_\infty} \\
&= \frac{\langle g, \psi \cdot \mathrm{dom}(g)\rangle - (\|\hat{\theta}\|_1 - \langle g, \psi \cdot \mathrm{dom}(g)\rangle) - \varepsilon\|\hat{\theta}\|_1}{p^{n^2} \cdot p^{-n^2}} \\
&= 2\langle g, \psi \cdot \mathrm{dom}(g)\rangle - (1+\varepsilon)\|\hat{\theta}\|_1 \\
&= 2 \cdot \left(1 + \frac{p - p^{-n+1}}{p-1}\right) p^{-n} \prod_{k=1}^{n} (p^k - 1) - (1+\varepsilon) \cdot p^{-n} \prod_{k=1}^{n} (p^k - 1) \prod_{k=0}^{n-1} \frac{1+p^k}{p^k} \\
&= \left(2\left(1 + \frac{p - p^{-n+1}}{p-1}\right) - (1+\varepsilon) \prod_{k=0}^{n-1} \frac{1+p^k}{p^k}\right) p^{-n} \prod_{k=1}^{n} (p^k - 1)
\end{aligned}$$

Note that

$$p^{-n} \prod_{k=1}^{n} (p^k - 1) \geq p^{-n} \prod_{k=1}^{n} p^{k-1} = p^{n(n-3)/2},$$

and when $\varepsilon = 1/4$,

$$2\left(1 + \frac{p - p^{-n+1}}{p-1}\right) - (1+\varepsilon) \prod_{k=0}^{n-1} \frac{1+p^k}{p^k} \geq 5.99 - (1+\varepsilon) \cdot 4.769 > 0.028, \qquad p = 2,\ n \geq 10,$$

$$2\left(1 + \frac{p - p^{-n+1}}{p-1}\right) - (1+\varepsilon) \prod_{k=0}^{n-1} \frac{1+p^k}{p^k} \geq 4 - (1+\varepsilon) \cdot 3.13 > 0.08, \qquad p \geq 3.$$

□

C Streaming

In all the linear algebra problems we have discussed, Alice and Bob want to know some property of the $n \times n$ matrix $x + y$, where Alice holds x and Bob holds y. However, if we want to compute the matrix property in the streaming model, the matrix in the stream is represented in row order. Hence in the communication model, we need a different way to distribute the inputs: Alice holds the top half of the $n \times n$ matrix and Bob holds the bottom half. Formally, Alice holds an $(n/2) \times n$ matrix x' and Bob holds another $(n/2) \times n$ matrix y'. They want to solve the problem on the concatenated matrix $\binom{x'}{y'}$. In some circumstances, this is a more natural way to distribute the input. We shall show that all of our lower bounds still hold even in this setting.

Theorem 29. *Alice and Bob hold $(n/2) \times n$ matrices x' and y' over $\mathbb{F}_p$, respectively. They need $\Omega(n^2 \log p)$ qubits of communication even if they use a quantum protocol to compute* $\textsc{Rank}'_{n,n-1}$, $\textsc{Inverse}'$ *or* $\textsc{LinSolve}'_b$.

- $\textsc{Rank}'_{n,n-1}$. *Decide if* $\operatorname{rank}\binom{x'}{y'}$ *is* n *or* $n-1$*;*
- $\textsc{Inverse}'$. *Decide if the top-left entry of* $\binom{x'}{y'}^{-1}$ *is zero;*
- $\textsc{LinSolve}'_b$. $b \in \mathbb{F}_p^n$ *is a non-zero vector.* $t \in \mathbb{F}_p^n$ *is the variable of the linear system* $\binom{x'}{y'}t = b$*. They want to decide if* t_0 *is zero.*

Proof. We reduce the original problems to these problems.

- $\textsc{Rank}_{n,n-1}$. Let $x' = \begin{pmatrix} x & -I \end{pmatrix}$, and $y' = \begin{pmatrix} y & I \end{pmatrix}$.

$$\operatorname{rank}(x+y) = \operatorname{rank}\begin{pmatrix} x+y & 0 \\ y & I \end{pmatrix} - n = \operatorname{rank}\begin{pmatrix} x & -I \\ y & I \end{pmatrix} - n = \operatorname{rank}\begin{pmatrix} x' \\ y' \end{pmatrix} - n.$$

  Therefore, $\operatorname{rank}(x+y) = n$ iff $\operatorname{rank}\binom{x'}{y'} = 2n$, and $\operatorname{rank}(x+y) = n-1$ iff $\operatorname{rank}\binom{x'}{y'} = 2n-1$.
- $\textsc{Inverse}'$. The reduction from $\textsc{Rank}'_{n,n-1}$ almost works. The only problem is that the parity of the size of the matrix changes after appending one row and one column. To make the size even, we add another row on the bottom and another column on the right. The additional row and column are all zero except for the bottom-right entry, and the bottom-right entry is one.
- $\textsc{LinSolve}'_b$. The reduction from $\textsc{Inverse}'$ still works. □

Finally, this implies the streaming part of Theorem 1.

D Privacy

We consider the $\textsc{rank}_{n,n-1}$ problem in this section. Let μ denote the uniform distribution over $G \times G$, where G is the semi-group considered in Section 3. Similarly to Theorem 5, we have

Theorem 30. *Suppose that* $\delta < 1/9$ *when* $p = 2$ *and* $\delta < \frac{p}{3(1-p+p^2)}$ *when* $p \geq 3$*. Then*

$$R^{pub}_{1/3}(\textsc{Rank}_{n,n-1}) \leq D^{\mu}_{\delta}(\textsc{rank}_{n,n-1}) \leq C\left(\log\frac{1}{\delta}\right) \cdot R^{pub}_{1/3}(\textsc{Rank}_{n,n-1}).$$

Proof. One can verify that

$$\alpha := \frac{\Pr\{\operatorname{rank}(A+B) = n\}}{\Pr\{\operatorname{rank}(A+B) = n-1\}} = \left(1 + \frac{1}{p^n - 1}\right)\frac{(p-1)^2}{p}.$$

Note that $\alpha \approx 1/2 < 1$ when $p = 2$ and $\alpha > 1$ when $p \geq 3$. Following the same reduction the same reduction in the proof of Theorem 5 with Π_ν replaced by Π_μ, we conclude that the public coin protocol Π' has error probability at most $1 - \frac{(1+\alpha)(1-\delta)-1}{\alpha} < \frac{1}{3}$ when $p = 2$ and at most $(1+\alpha)\delta < \frac{1}{3}$ when $p \geq 3$. The rest follows similarly as in Theorem 5. □

As a corollary of [4, Theorem 1.3], we know that when p is a constant,

$$\mathrm{IC}_\mu(\textsc{Singularity}_n) \geq \mathrm{IC}_\mu(\textsc{Rank}_{n,n-1}) = \Omega\left(\frac{D^\mu(\textsc{Rank}_{n,n-1})}{\operatorname{polylog} D^\mu(\textsc{Rank}_{n,n-1})}\right) = \Omega\left(\frac{n^2}{\operatorname{polylog} n}\right).$$

We remark that combining [13] and [15] yields $\max_\lambda \mathrm{IC}_\lambda(\textsc{Rank}_{n,n-1}) = \Omega(n^2)$, but it is not clear what distribution λ attains the lower bound. Our bound above, although slightly weaker, shows that the product distribution nearly achieves the desired lower bound. Finally, it then follows from [1, Proposition 20] that

$$\mathrm{PRIV}_\mu(\textsc{Singularity}_n) = \Omega\left(\frac{n^2}{\operatorname{polylog} n}\right).$$